

ABSTRACT OF THE DISCLOSURE

A computer architecture includes a first operating system (COS), which may be a commodity operating system, and a kernel, which acts as a second operating system. The COS is used to boot the system as a whole. After booting, the kernel is loaded and displaces the COS from the system level, meaning that the kernel itself directly accesses predetermined physical resources of the computer. All requests for use of system resources then pass via the kernel. System resources are divided into those that, in order to maximize speed, are controlled exclusively by the kernel, those that the kernel allows the COS to handle exclusively, and those for which control is shared by the kernel and COS. In the preferred embodiment of the invention, at least one virtual machine (VM) runs via a virtual machine monitor, which is installed to run on the kernel. Each VM, the COS, and even each processor in a multiprocessor embodiment, are treated as separately schedulable entities that are scheduled by the kernel. Mechanisms for high-speed I/O between VM's and I/O devices are also included.